University of California, Berkeley
College of Engineering
Computer Science Division — EECS

Spring 1998 D.A. Patterson

Quiz 1
March 4, 1998
CS252 Graduate Computer Architecture

You are allowed to use a calculator and one 8.5" by 11" double-sided page of notes. Show your work on all problems. Good luck!

Your name:	
SID number:	
Email address:	

1	/12
2	/20
3	/24
Total	**/56**

Question 1: Calculate Your Cache

A certain system with a 350 MHz clock uses a separate data and instruction cache, and a unified second-level cache. The first-level data cache is a direct-mapped, write-through, write-allocate cache with 8kBytes of data total and 8-Byte blocks, and has a perfect write buffer (never causes any stalls). The first-level instruction cache is a direct-mapped cache with 4kBytes of data total and 8-Byte blocks. The second-level cache is a two-way set associative, write-back, write-allocate cache with 2MBytes of data total and 32-Byte blocks.

The first-level instruction cache has a miss rate of 2%. The first-level data cache has a miss rate of 15%. The unified second-level cache has a local miss rate of 10%. Assume that 40% of all instructions are data memory accesses; 60% of those are loads, and 40% are stores. Assume that 50% of the blocks in the second-level cache are dirty at any time. Assume that there is no optimization for fast reads on an L1 or L2 cache miss.

All first-level cache hits cause no stalls. The second-level hit time is 10 cycles. (That means that the L1 miss penalty, assuming a hit in the L2 cache, is 10 cycles.) Main memory access time is 100 cycles to the first bus width of data; after that, the memory system can deliver consecutive bus widths of data on each following cycle. Outstanding non-consecutive memory requests can not overlap; an access to one memory location must complete before an access to another memory location can begin. There is a 128-bit bus from memory to the L2 cache, and a 64-bit bus from both L1 caches to the L2 cache. Assume a perfect TLB for this problem (never causes any stalls).

a) What percent of all data memory references cause a main memory access (main memory is accessed before the memory request is satisfied)? **First show the equation, then the numeric result.**

b) How many bits are used to index each of the caches? Assume the caches are presented physical addresses.

Question 1 (continued)

c) How many cycles can the longest possible data memory access take? Describe (briefly) the events that occur during this access.

d) What is the average memory access time in cycles (including instruction and data memory references)? **First show the equation, then the numeric result.**

Question 2: Tomasulo's Revenge

Using the DLX code shown below, show the state of the Reservation stations, Reorder buffers, and FP register status for a speculative processor implementing Tomasulo's algorithm as presented in the book. Assume the following:

- Only one instruction can issue per cycle.
- The reorder buffer has 8 slots.
- The reorder buffer implements the functionality of the load buffers and store buffers.
- All FUs are fully pipelined.
- There are 2 FP multiply reservation stations.
- There are 3 FP add reservation stations.
- There are 3 integer reservation stations, which also execute load and store instructions.
- No exceptions occur during the execution of this code.
- All integer operations require 1 execution cycle. Memory requests occur and complete in this cycle. (For this problem, assume that, barring structural hazards, loads issue in one cycle, execute in the next, write in the third, and a dependent instruction can start execution on the fourth.)
- All FP multiply operations require 4 execution cycles.
- All FP addition operations require 2 execution cycles.
- On a CDB write conflict, the instruction issued earlier gets priority.
- Execution for a dependent instruction can begin on the cycle after its operand is broadcast on the CDB.
- If any item changes from "Busy" to "Not Busy", you should update the "Busy" column to reflect this, but you should not erase any other information in the row (unless another instruction then overwrites that information).
- Assume the all reservation stations, reorder buffers, and functional units were empty and not busy when the code show below began execution.
- The "Value" column gets updated when the value is broadcast on the CDB.

Integer registers are not shown, and you do not have to show their state.

Question 2 (continued)

a) The tables below show the state after the cycle in which the second `SUBI` from the code below issued. Show the state after the next cycle.

```
Lp:  LD     F0, 0(R1)
     LD     F2, 0(R2)
     MULTD  F4, F0, F2
     ADDD   F6, F0, F0
     SUBI   R1, R1, 8
     SUBI   R2, R2, 8
     ADDI   R3, R3, 1
```

Reservation stations							
Name	Busy	Op	Vj	Vk	Qj	Qk	Dest
Add1	Y	`ADDD`	F0	F0			#4
Add2							
Add3							
Mult1	Y	`MULTD`	F0	F2			#3
Mult2							
Int1	Y	`SUBI`	R2	8			#6
Int2	N	`LD`	R2	0			#2
Int3	Y	`SUBI`	R1	8			#5

Reorder buffer					
Entry	Busy	Instruction	State	Destination	Value
1	N	`LD F0, 0(R1)`	Commit	F0	Mem[0(R1)]
2	N	`LD F2, 0(R2)`	Commit	F2	Mem[0(R2)]
3	Y	`MULTD F4, F0, F2`	Execute	F4	
4	Y	`ADDD F6, F0, F0`	Execute	F6	
5	Y	`SUBI R1, R1, 8`	Execute	R1	
6	Y	`SUBI R2, R2, 8`	Issue	R2	
7					
8					

FP register status									
Field	F0	F2	F4	F6	F8	F10	F12	...	F30
Reorder #	1	2	3	4				...	
Busy	N	N	Y	Y				...	

Question 2 (continued)

b) The tables below show the state during the cycle in which the second `MULTD` from the code below issued. Show the state after the next ***two*** cycles.

```
MULTD  F0, F2, F4
ADDD   F6, F6, F0
ADDD   F2, F2, F8
LD     F4, 0(R2)
ADDI   R1, R1, 8
MULTD  F8, F10, F12
ADDD   F4, F4, F10
ADDI   R2, R2, 1
```

Reservation stations							
Name	Busy	Op	Vj	Vk	Qj	Qk	Dest
Add1	Y	`ADDD`	F6			#1	#2
Add2	Y	`ADDD`	F2	F8			#3
Add3							
Mult1	N	`MULTD`	F2	F4			#1
Mult2	Y	`MULTD`	F10	F12			#6
Int1	Y	`LD`	R2	0			#4
Int2	Y	`ADDI`	R1	8			#5
Int3							

Reorder buffer					
Entry	Busy	Instruction	State	Destination	Value
1	Y	`MULTD F0, F2, F4`	Write	F0	F2×F4
2	Y	`ADDD F6, F6, F0`	Issue	F6	
3	Y	`ADDD F2, F2, F8`	Execute	F2	
4	Y	`LD F4, 0(R2)`	Execute	F4	
5	Y	`ADDI R1, R1, 8`	Execute	R1	
6	Y	`MULTD F8, F10, F12`	Issue	F8	
7					
8					

FP register status									
Field	F0	F2	F4	F6	F8	F10	F12	...	F30
Reorder #	1	3	4	2	6			...	
Busy	N	Y	Y	Y	Y			...	

Question 2 (continued)

c) The tables below show the state after the cycle in which the `SUBD` from the code below issued. Show the state after the next ***four*** cycles.

```
ADDD   F4, F0, F0
SUBD   F4, F4, F2
ADDI   R2, R2, 1
ADDI   R3, R3, 1
ADDD   F2, F6, F8
MULTD  F0, F6, F8
```

Reservation stations							
Name	Busy	Op	Vj	Vk	Qj	Qk	Dest
Add1	Y	ADDD	F0	F0			#1
Add2	Y	SUBD		F2	#1		#2
Add3							
Mult1							
Mult2							
Int1							
Int2							
Int3							

Reorder buffer					
Entry	Busy	Instruction	State	Destination	Value
1	Y	`ADDD F4, F0, F0`	Execute	F4	
2	Y	`SUBD F4, F4, F2`	Issue	F4	
3					
4					
5					
6					
7					
8					

FP register status									
Field	F0	F2	F4	F6	F8	F10	F12	...	F30
Reorder #			2					...	
Busy			Y					...	

Question 3: Vector vs DSP Showdown

Examine the two architectures below.

The first architecture is a 25 Mhz 3-stage dsp processor. A block diagram showing some of the fully-bypassed datapath is shown below. The three stages are fetch, decode (where branches

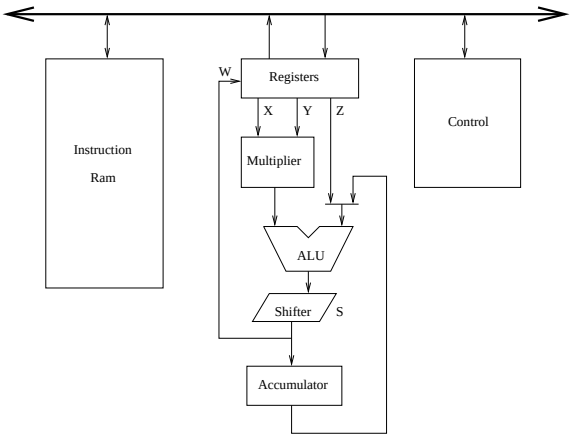

Figure 1: The DSP block diagram

are evaluated and the PC updated), and execute (where memory and register writes also occur). The processor is able to multiply, accumulate, and shift during its execute stage. It has the same load, store, and branch instructions as DLX. It also includes a `LT` instruction, which loads a value into a register from memory, and decrements the base register to the next element. The arithmetic operations are slightly different:

- Register 0 always contains the value 0
- Register 1 always contains the value 1
- The result from the shifter is always written to the accumulator on arithmetic operations
- Operations can be specified as `MAC W, X, Y, Z, S`, where `W` is the register to be written; `X` and `Y` are registers that go to the multiplier; `Z` is the register that goes to the alu; and `S` specifies the amount to right shift the result.
- Operations can also be specified as `MACA W, X, Y, S`, where `W` is the register to be written; `X` and `Y` are registers that go to the multiplier; the accumulator goes to the alu; and `S` specifies the amount to right shift the result.

The second architecture is a 100 Mhz vector processor with a MVL of 64 elements. It has one FP add/subtract FU, one FP multiply/divide FU, and a single memory FU. The startup overhead is 5 cycle for add, subtract, multiply, and divide instructions, and 10 cycles for memory instructions. It supports flexible chaining but not tailgating.

Question 3 (continued)

Here is the code for the dsp:

```
LP: LT     R2, 0(R5)          # Load R2 with new value
    MAC    R0, R10, R2, R0, 0 # Perform the calculation
    MACA   R0, R11, R2, 0
    MACA   R2, R12, R2, 0
    BNEZ   R5, LP
    SW     R2, 4(R5)          # Delayed branch
```

a) What is the peak performance, in results per second, of the above three-tap filter? Count each tap as a result.

b) What would be the peak performance, in results per second, of the above code, if it was a five-tap filter? Count each tap as a result.

Question 3 (continued)

c) Translate the following DLX code to code that will operate on this DSP. Assume that all floating point calculations below can be done in fixed-point on the DSP. Do not worry about round-off error from converting between floating point and fixed point. Assume that for DLX code, F0 has been initialized to 0, F2 to 0.5, and F4 to 2.0. Assume that for the DSP code you will write, R2 has been initialized to 2. Assume for both that register R5 contains the correct initial loop count.

```
LP: LW      R3, 0(R5)    # Load R3 with the new value
    MOVI2FP F6, R3       # Move the value from R3 to F6
    CVTI2F  F6, F6       # Convert integer value to floating point
    MULTF   F8, F4, F6   # Multiply by 2
    ADDD    F10, F8, F0  # Add accumulator to value
    MULTF   F0, F10, F2  # Divide value by 2
    CVTF2I  F12, F0      # Convert to integer representation
    MOVFP2I R3, F12      # Move it to integer registers
    SW      R3, 0(R5)    # Store value back to mem
    ADDI    R5, -4       # Point to next element
    BNEZ    R5, LP       # If not done, branch back
```

Question 3 (continued)

Here is the code for the vector machine:

```
LP: LV     V1, R5           # Load V1 with new value
    MULTSV V2, F0, V1       # Perform the calculation
    MULTSV V3, F1, V1
    ADDV   V2, V2, V3
    MULTSV V3, F2, V1
    ADDV   V2, V2, V3
    SV     V2, R5
    SUBI   R5, R5, 8
    BNEZ   R5, LP
```

d) Show the convoys of vector instructions for the above code. Follow the timing examples in the book. Draw lines to show the convoys on the existing code shown above.

e) Show the the execution time in clock cycles of this loop with n elements (T_n); assume T_{loop} = 15. **Show the equation, and give the numeric value of execution time for n=64.**

f) What is R_∞ for this loop? **Show the equation, and give the numeric value.**

Question 3 (continued)

g) List 6 characteristics of DSP instruction set architectures that differ from general purpose microprocessors.

h) Which of those characteristics are supported in vector architectures?

i) Which of the unsupported characteristics could be handled in software?

j) What changes to the hardware would you make to handle the remaining characteristics?